Case Study

Reducing Ingestion Latency from 8 Hours to Under 10 Minutes Across a 420+ Branch Network

Client Overview

Large-scale financial and insurance networks manage vast, distributed branch infrastructures where data sync delays directly limit operational agility. This case study details the deployment of a real-time, log-based Change Data Capture (CDC) ingestion framework for a leading insurance provider operating 420+ locations with a revenue base of over 71,000 Crores. By implementing a centralized, agentless data replication platform, the project successfully modernized legacy 6–8 hour nightly batch processing delays into a continuous stream that delivers data across hundreds of endpoints in under 10 minutes.

The Challenge?

Batch Ingestion Latency

High operational dependency on overnight batch pipelines requiring 6–8 hours to execute, which delayed morning dashboard availability.

Endpoint Proliferation

Inability to efficiently migrate, consolidate, and synchronize diverse data formats across a sprawling hybrid network of hundreds of targets.

Source Degradation Risks

Severe processing friction and resource competition caused by traditional, query-heavy data replication tools hitting active transactional mainframes.

Solution Delivered

Agentless Replication Architecture

High-throughput pipelines equipped with native data pseudonymization modules to securely ingest data while ensuring strict compliance with public sector privacy regulations.

Log-Based CDC Automation

Machine learning forecasting algorithms built to model multi-decade demographic and socioeconomic variables using Python and R.

Heterogeneous Pipeline Translation

A robust, enterprise-grade Azure SQL data asset combining scattered public datasets into a centralized query repository.

Resilient High-Availability Infrastructure

Tailored Power BI and Tableau reporting frameworks serving aggregate regional trends to council executives and targeted risk indicators to frontline caseworkers.

Architecture

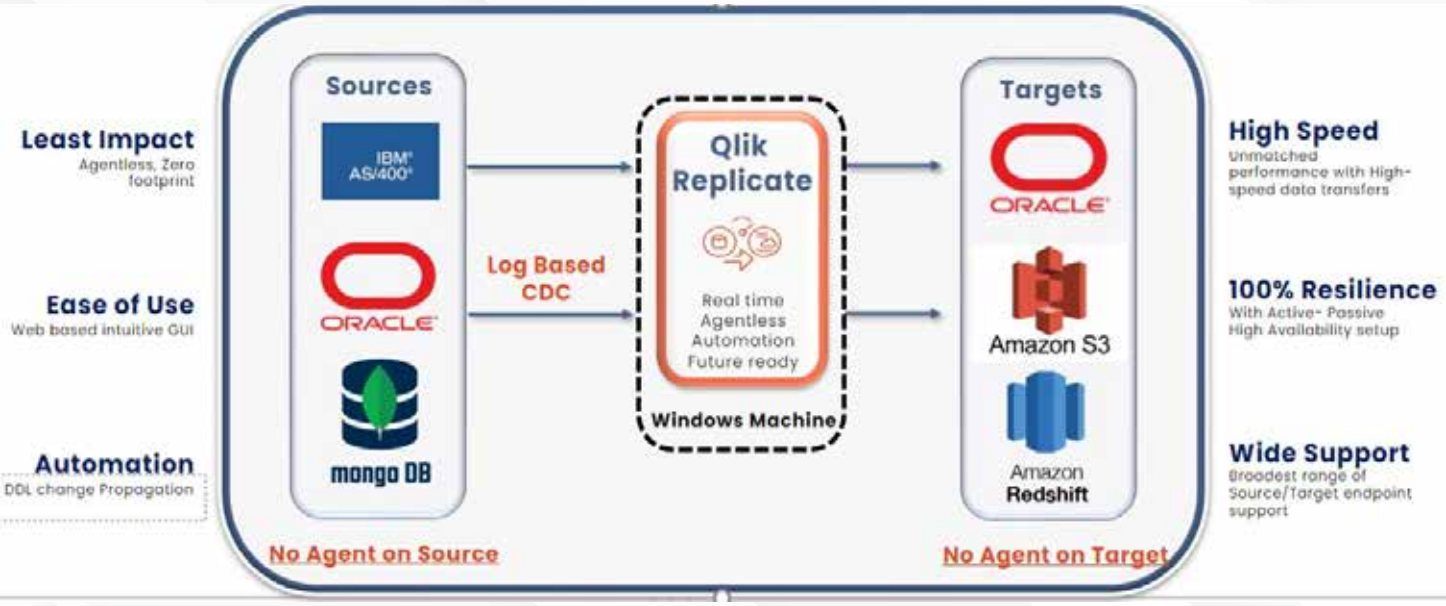

Impact

Substantial Performance Gains

Replaced time-consuming nightly processing to achieve guaranteed data freshness within an under 10-minute synchronization window.

Touchless Operational Model

Achieved 100% manual overhead reduction by creating automated tasks that handle initial mass loads and switch to active tracking automatically.

Preserved Core Performance

Safeguarded critical front-line production database performance by moving integration tasks to low-impact log scans

Standardized Infrastructure Control

Consolidated multiple disconnected ingestion scripts into an intuitive graphical system capable of managing hundreds of endpoints.

Outcome

Transitioning away from slow, batch-heavy processing to a modern stream architecture is essential to staying competitive in high-volume financial markets. By uniting low-impact log-based CDC with an intuitive, agentless replication framework, this prominent insurance carrier successfully transformed its enterprise data velocity. The solution delivers the processing speed, endpoint flexibility, and absolute architectural resilience required to lower overhead, satisfy operational SLAs, and support an agile, data-first corporate culture at scale.

Interested in a customized solution designed to meet your unique requirements and achieve comparable results?

At Exponentia.ai, we partner with businesses to address complex data challenges and build trust in enterprise data. Our approach helps organizations streamline data management, improve visibility, and empower teams with reliable insights for better decision-making.

Engage with Us

www.exponentia.ai

engage@exponentia.ai

US UK UAE IND SGP